

Applying a Testing Methodology to Augmented Reality Interfaces to Simulation Systems

Mark A. Livingston*

Dennis Brown*

J. Edward Swan II[†]

Brian Goldiez[‡]

Yohan Baillet[§]

Greg S. Schmidt[§]

Naval Research Laboratory
Washington D.C.

Abstract

Mobile augmented reality (AR) combines 3D spatially registered graphics and sounds with a user's perception of the real world. Combining mobile AR with computer simulation promises to revolutionize practicing and training for many tasks, especially those that are naturally conducted in the real world. However, to achieve this potential, the field needs a much better understanding of how users perceive and comprehend information that is mediated by an AR display. We review the work that has been performed in this area to date and discuss the challenges presented by perceptual issues in AR for training systems. Then we describe our application of experimental methodologies which address perceptual and cognitive issues, employs subject matter experts to ensure domain relevance, and address the limitations of emerging technology. We apply these methodologies in the context of a mobile AR simulation system that we have developed to support military training.

1 Introduction

Virtual simulations of military tasks are generally accepted as a useful method of training for the U.S. military [32, 7]. What is unknown, however, is the effect of the visual fidelity and behavioral realism of the simulation on the utility of the system. These factors are difficult to provide in simulations.

In an effort to create a training method that combines the control and repeatability of VR with the authenticity of the real world, we have implemented a prototype training system using *augmented reality* (AR), a display technique which mixes virtual cues with the user's perception of a real environment. Figure 1 shows the real world (through a camera) augmented by a set of virtual targets (vehicles and tanks). The same paradigm can be used to display battlefield intelligence information to a dismounted warfighter in a head-up manner, similar to the head-up display systems for pilots.

A mobile AR system contains all components necessary to display AR in a self-contained and portable package. Figure 2 shows the mobile AR system developed for the Battlefield Augmented Reality System (BARS) [18]. The system tracks the user's position and orientation (using a GPS and inertial sensors). This information is fed to a wearable computer (PC-compatible laptop or board computer) which generates 3D graphics. These graphics are displayed on a head-



Figure 1: A sample view of our system, showing virtual vehicles and smoke within a virtual battlespace in a real training facility.

mounted display (Sony Glasstron, Microvision Nomad, or Trivision). This approach integrates spatial information with objects in the environment.

For training, animated 3D computer-generated forces are inserted into the environment. The AR training system moves the repeatability and control of a VR system into a real-world training environment. We refer to this variation of the systems as BARS-Embedded Trainer (BARS-ET) [3, 4, 5]. Existing training facilities can be used with BARS-ET, but with training scenarios that are limited only by the power of the computer simulations. Animated computer-generated forces (CGFs) appear on the display, properly registered and occluded in the real world. The CGF behaviors are controlled by a Semi-Automated Forces (SAF) system.

Such a system raises a number of perceptual and cognitive issues arising from the question of how the fidelity of the synthetically-generated cues affects training effectiveness. Researchers have been addressing the perceptual and cognitive issues in AR from two perspectives. *Low-level* tasks develop understanding of how human perception and cognition operate in AR contexts. *High-level* applications show how AR could impact underlying tasks, leveraging domain analysis with subject matter experts. These approaches are complimentary, and our team has found success by integrating them [14, 18, 17] in the development of a mobile AR simulation system to support military training.

We believe that combining the approaches into a single methodology will enable us to evaluate both system capabilities and user performance with the system. Domain analysis will help ensure that the system includes the most useful ca-

*Virtual Reality Laboratory, Naval Research Laboratory. Corresponding email: mark.livingston@nrl.navy.mil

[†]Current address: Dept. of Computer Science, Mississippi State Univ.

[‡]University of Central Florida

[§]ITT Advanced Engineering and Sciences / Naval Research Laboratory

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2005		2. REPORT TYPE		3. DATES COVERED 00-00-2005 to 00-00-2005	
4. TITLE AND SUBTITLE Applying a Testing Methodology to Augmented Reality Interfaces to Simulation Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Washington, DC, 20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Western Simulation Multiconference (WMC '05), New Orleans, LA, January 23-27, 2005.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			



Figure 2: A user wearing our system. The inset right shows a generic view of virtual spatial data aligned with the real environment.

pabilities. Human perception is an innate ability, and variations in performance will reflect the system’s appropriateness for use as a training tool. Thus, we are really evaluating the system’s performance by measuring the user’s performance on perceptual-level tasks. The evaluation of cognitive-level tasks will enable us to determine how users are performing. Such high-level metrics can only be measured after the results of the perceptual-level tests inform the system design.

2 Related Work

2.1 Perceptual tests

Understanding how a user will perceive the presented data is a first step in learning how a visualization system helps (or fails to help) the user. One important question for our system, as discussed in Section 4.1 is *the perception of depth*. A number of studies have shown representations that convey depth relationships between real and virtual objects. Partial transparency, dashed lines, overlays, and virtual cut-away views all give the user the impression of a difference in the depth [10, 27, 31, 21]. These studies use simple tasks, but can thus focus on the representation.

A study of the presence of a visible (real) surface near a virtual object found a significant influence on the user’s perception of the depth of the virtual object in the near field [9]. For most users, the virtual object appeared to be nearer than it really was. This varied widely with the user’s age and ability to use accommodation, even to the point of some users being influenced to think that the virtual object was further away than it really was. Adding virtual backgrounds with texture reduced the errors, as did the introduction of virtual holes, similar to those described above. Further experimentation showed that superposition of the real and virtual was needed to produce this change; a surface with a hole did not produce such a change in the perception of the virtual object. Latency linearly degraded the user’s ability to match the location of a physical pointer to a virtual object [19]. Occlusion of the real object by the

virtual object gives the incorrect impression that the virtual object was in front, despite the object being located behind the real object and other perceptual cues denoting this relationship [25]. Further studies showed that users performed better when allowed to adjust the depth of virtual objects than when making forced-choice decisions about the objects’ locations [24].

A pilot study using video AR [13] showed users a stimulus which was either behind or at the same distance as an occluding surface. The users identified whether the stimulus was behind, at the same distance as, or closer than the occluder. The performance metric is thus an ordinal depth measure. Only a single occluded object was present in the test. The parameters in the pilot test were the presence of a cutaway in the obstruction and motion parallax. The presence of the cutaway significantly improved users’ perceptions of the correct location when the stimulus was behind the obstruction. The authors offered three possible locations to the users, even though only two locations were used. Users consistently believed that the stimulus was in front of the obstruction, despite the fact that it was never there.

Our research uses representational cues such as drawing style and opacity to improve the user’s perception of the depth relationships of multiple occluded objects [17]. Initial work has found that user error in identifying ordinal depth relationships was lower with a drawing style that used a filled object with a wireframe outline and decreasing opacity with distance when compared with the use of a consistent ground plane constraint. As a secondary result, a positive main effect of repetition on response time but not on accuracy indicated that subjects quickly understood the semantic meaning of the encodings.

Shadows give important depth cues, and their inclusion in a personal-space AR representation improved the user’s ability to identify proper depth ordering, but did not improve ordering within the fronto-parallel plane [29]. Users in this study on shadows did not employ motion parallax except in the condition of monocular view and low-angle vantage point, which is not surprising since all objects were in the near field. Users reported a subjective rating that shadows made objects more aesthetically pleasing.

2.2 Task-oriented Studies

One example of a task-oriented study is the application of AR to medical interventions with ultrasound guidance [26]. A doctor performed ultrasound-guided needle biopsies with and without the assistance of an AR system designed for the task. A second physician evaluated the needle placement. Needle localization improved when using the AR system. The performance metric was the standard for evaluating doctors’ performance: needle placement at a pre-specified set of locations within the target lesion. The physician uses the ultrasound to determine the ideal and actual needle locations. Thus the measure is tightly connected to the task, and in fact is independent of the AR system.

Two comparisons have been done for mechanical tasks with and without AR [23, 30]. The studies found that the time required to complete the task was significantly less with AR than with only a printed manual. This was largely due to a reduced need for switching context between the task area and the instructions. The AR system was able to embed instructions within the environment. One study measured a reduced user mental workload and error rate [30]. An earlier test found no significant difference in the time required and attributed this to difficulties with the user interface [6].



Figure 3: *Left:* Optical see-through leaves synthetic objects with a ghostly appearance. *Right:* Graphical overlay on video enables full occlusion of real objects by synthetic objects.

These two tasks represent the only interactive applications of AR technology that have been successful to date. The mechanical assembly tasks represent a limited implementation of AR, and the medical applications allow the system designers to exert considerable control over the surrounding environment. This has limited the use of testing methodologies for such systems.

3 Perceptual Issues in AR for Training and Operations

There are two primary sensory channels we use with our AR system, visual and aural. We consider three sources of conflict in perceptual cues: properties of the display, rendering parameters, and audio presentation.

3.1 Properties of HMD

There are two fundamentally different technologies used for AR graphic displays: optical see-through and video see-through¹. The type of head-mounted display used can make or break the illusion of CGFs existing in the real world. If the graphical representations do not occlude the appropriate elements of real world, the virtual objects do not appear solid or realistic. Instead, the graphics, translucent on a non-occluding display, take on a ghostly and unrealistic appearance. Figure 3 shows two similar scenes, one through a non-occluding display, and one through an occluding display. Notice how the avatar is washed out in bright light in the non-occluding display. It is clear that this will significantly impact the user's experience of the system. However, it remains to be seen whether such low-fidelity sensory cues adversely impact the effectiveness of the training received in such a system.

The video-based display accumulates latency waiting for the camera to capture and transfer an image of the real environment, thus showing the image later than an unadorned user would perceive it. This effect is noticeable by users but thus far appears to be less problematic than other system errors or otherwise tolerated. This may be because users were often dealing with (nearly) static environments and can thus wait for the system to catch up with reality.

Most video-based displays implemented to date have an offset between the camera's apparent location and the user's eyes. Clearly, a physical offset is necessary, but a cleverly-designed optical path using mirrors can reduce or eliminate the apparent offset [12]. If this is not done, the offset interferes with the user's hand-eye coordination during use of

the video-based AR display, and the adaptation of the user back to "normal" hand-eye coordination is slow [1].

For mobile AR for situation awareness (not training), the optical display, even with its faults, is the better of the two types of displays because the user's view of the real world is not degraded and the ghostly appearance of tactical information does not detract from the utility of that information. For embedded training, however, the benefits of the video display's complete occlusion of the real objects by virtual objects outweigh the drawback of decreased resolution. So it is our choice until an optical see-through display with complete occlusion capabilities [16] becomes widely available.

3.2 Rendering Parameters

The first implementation of BARS-ET used static VRML models for the computer-generated forces, and seeing the static models slide around the environment was not convincing to the first users of the system. Adding realistically animated humans to the system was another low-impact improvement that paid off well. In this case, only a third-party software library was added. The DI-Guy animation system [2] was integrated into the BARS-ET graphics renderer. Combined with the occlusion model, the forces realistically emerge from buildings and walk around corners.

Model fidelity is controlled by the modeler and is limited by the power of the machine running the application. Although models that can be rendered in real time still look computer-generated, just like in VR-based simulations, the limited model representation capabilities are adequately realistic for embedded simulation and training. AR actually has an advantage over VR with respect to rendering: the AR graphics system does not need to draw an entire virtual world, only the augmented forces, so they could potentially be more detailed than those in VR-based simulations.

Lighting virtual objects is a problem we have not approached yet. It would require knowing the lighting conditions of the real environment in which the model would appear, and changing the renderer's light model to match. Another limitation is the display itself, as it is very sensitive to outside light. Even if the image is rendered with perfect lighting, it still might not appear correct to the user.

Occlusion of distant objects by close objects is a powerful depth cue. We would like to have nearby virtual objects occlude distant real objects. To provide this, we use the viewpoint and a model of the real environment, both available in the AR system for situation awareness. In that system, this data enables graphical cues to be properly located and scaled on the display and even provide the user with information about occluded objects. For embedded training, we perform a depth comparison of the virtual objects against the real world in order to show only the closer object (real or virtual) at each pixel that a virtual object might occupy. This solution was introduced for indoor applications by [28] and applied to outdoor models by [22] for use in outdoor AR gaming. The utility of allowing a trainee to have the ability to see virtual objects through occluding surfaces is at best premature before this ability exists in military equipment, and at worst detrimental by making the training too easy. We have yet to perform experiments to examine this question.

In the use of BARS-ET for training for operations, the user carries a simulated weapon. Tracking such a hand-held device is a classic problem for an AR system. While there are some good indoor systems, tracking outdoors remains elusive.

¹Rolland and Fuchs [25] discuss advantages and disadvantages for medical applications, many of which apply to other applications.

3.3 Inserting Aural Cues

Spatialized audio enhances the training experience by giving the user audio information that matches the visual display of the environment (e.g. footsteps of virtual soldiers) and making the experience more realistic and memorable [8]. To render the graphical display, the system tracks the user's attitude in the virtual world along with locations of simulated military forces. This data also supports spatialized audio. Sounds can be attached to virtual objects (e.g. helicopters) or events (e.g. gunfire). A 3D sound API is updated continuously with the positions of the user and simulated entities. BARS-ET supports the Virtual Audio Server [11] and Microsoft's DirectX [20]. The API takes simple monophonic sound files and renders them in the user's headphones so that they sound like they have distinct positions in the real world. Open-air headphones naturally mix the sounds of the real world with the computer-generated sounds. The audio features have yet to be included in any user studies, however.

4 Application of Experimental Methodology

As AR technology begins to mature, we and some other research groups are considering how to test user perception and cognition when aided by AR systems. We give two examples of our application of such methodology. One focuses more on the perceptual level, while one focuses more on the task-performance level.

4.1 Depth Matching for Targeting

We determined a task for evaluation of BARS that was appropriate for operational use by dismounted warfighters [14]. For BARS-ET, however, we need a slightly different task. The task for BARS was to judge depth of occluded objects through virtual representations. This is a physically unrealistic scenario, which is inappropriate for training.

For BARS-ET, we need to modify the task to test perception of depth of virtual representations of objects that would be physically visible if they were real. This is important because an envisioned task for BARS-ET users is that of a forward observer, who call coordinates for indirect fire on a target. The forward observer thus maintains line-of-sight contact with the target, and determines the direction and distance from his location to the target.

Results from our pilot test (Figure 4) showed that users were reasonably accurate at matching distances, but got poorer with increasing distance to the real object (Figure 5). One training scenario has the forward observer look at a miniature world of about 60 m; our hallway scene measured about 45 m to the most distant referent, with referents roughly equi-spaced along the usable length of the hallway.

Four users (normal or corrected-to-normal) manipulated the virtual target's depth in the hallway with a trackball that was restricted to one degree of freedom in software. Linear perspective, provided by the walls and apparent size of the referents and virtual target, was the most powerful depth cue available. The virtual target also became more transparent as the user moved it further away ($\alpha \in [0.7, 0.3]$, near to far).

The next step for this particular task is to design a study that can measure the effectiveness of the training with the virtual targets. As noted, virtual targets enable capabilities that are difficult to attain with current targets, such as mobility and variety of scenarios. But because the nature of the depth perception task is not yet fully understood, we investigate the perceptual effects of this type of training.



Figure 4: User's view of the depth-matching task. The user matches the virtual target to the referent of the same color (here, green).




Figure 5: Depth-matching accuracy in a pilot study. The red line indicates the positions of the real referents. The black line indicates the mean virtual position of the virtual target within the same space. The boxes indicate standard deviation; asteriks indicate outliers.

The appropriate measures for a study of the overall effectiveness are two-fold. First, there are objective measures such as the number of repetitions required to reach a desired level of proficiency. We also believe that subjective measures such as user confidence in their performance and frustration with the system will be relevant.

The design for such a study highlights another aspect of our methodology, that of choosing appropriate users to be subjects. For perceptual studies, any user with normal or corrected-to-normal perception is appropriate. For task-based experiments, we expect our end users to be active military personnel or have military experience. As the task becomes more specialized, so must our users.

4.2 Navigation

AR offers the possibility to support navigational tasks in training or operational settings. We are investigating a special case called search and rescue navigation. This involves




Figure 6: A schematic of the maze used in the navigation study. The red dot indicates the location of the search target.

seeking an objective (e.g. a hostage) and completely traversing an unfamiliar space (e.g. to neutralize hostile situations).

The studies [15] use BARS with a $16^\circ \times 20^\circ$ monocular display which contains a map of an area. This area is a small maze (Figure 6) that is modeled in BARS and in a physical maze of 15×15 feet. This level of AR represents augmentation that is added to the real world, not integrated (or aligned) like the virtual targets.

The independent variable of greatest interest was the type of map. Types used were: a self-orienting virtual map that turned as the user turned, a static virtual map, and a paper map. Additionally, some users were given the ability to control the presence of the virtual map types. This study used 120 novices, balanced for gender. Preliminary results show that maze coverage was improved with an on-demand map view, but that people were fastest with paper maps. There was no significant effect of the type of map on the users' ability to find objects. Users were shown several overhead-view maps of a maze after they completed the exercise and asked which maze represented the one they had just completed. There was no significant effect of the type of map on their ability to recognize the correct maze.

Perceptual issues, such as layout and legibility of annotations, affect the user's ability to make use of the AR navigation assistance. We collaborate with researchers in these areas; however, experimental data on the perceptual components are not yet available. We did not want to allow the technology limitations to delay the test from being run. But we did control background illumination to maximize legibility, for example. Based on the preliminary results of the study, perceptual tests of layout and color will have much data with which to inform further designs of the maze study.

5 Discussion and Conclusions

Testing the effectiveness of AR-based training systems seems like a natural step in evaluating the benefits of such systems. Clearly, there are features available in such systems that are not available in current training methods, so there is bound to be utility in AR for training. But in order to be sure that there are no negative training effects, we must evaluate the perceptual and cognitive processes and ensure that they are consistent with the task on which the user is being trained. Were the system to cause the user to focus on the interface

technology, the training benefit would be adversely affected. At the same time, we sometimes create constrained scenarios in order to reduce the effects of such confounding factors.

We believe training instructors should be able to exert more direct control over the AR-based training system. Rather than coaching a team of actors in a training scenario, an instructor could interactively control synthetic forces, offering variety in training or the ability to repeat scenarios. The AR-based trainer also reduces costs (fewer people, virtual destruction of targets) and time required to re-initialize scenarios. The ultimate benefit to the end user of testing such systems is the increased training effect that should come from improvements to the system.

We argue for a progression from perceptual to cognitive tests. This is partly to allow the system to improve rather than being tested in a sub-standard state. But it also enables us to ensure that the basic system functions are well-suited to the users' needs before the variables that we can measure are intertwined into factors that can not be separated. For example, by studying the depth representations such as geometric depiction, shadows, and similar perceptual factors, we know that an observer's ability to identify coordinates of a target is not limited due to misleading depth information conveyed by the system. While we can not make absolute guarantees that a perfect system will be found, we can measure the effectiveness of the individual components and identify which system components we must improve to be most helpful to the user.

We have noticed some factors for future studies in performing the experiments we have described. Some users could not adjust for the deficiencies of the video-based display and felt they were in a completely virtual world, which works against the purpose of an AR training system. Most users did not experience this problem. Virtual objects should give the illusion that they exist in the real world and behave by its rules. There are several inherent problems: fidelity or realism in both static and dynamic models, lighting to match the real environment, and occlusion by real objects. The level of fidelity required for effective training is a subject for future work.

We believe AR will be an effective training tool for many situations, just as virtual environments have proven successful. Through a rigorous process of testing what features are beneficial and what features interfere with training, we hope to bring AR systems to a level which will allow the military to integrate them into effective training tools.

References

- [1] Frank A. Biocca and Jannick P. Rolland. Virtual eyes can rearrange your body: Adaptation to visual displacement in see-through, head-mounted displays. *Presence: Teleoperators and Virtual Environments*, 7(3):262–277, June 1998.
- [2] Boston Dynamics. DI Guy, October 2004. <http://www.bdi.com/content/sec.php?section=diguy>.
- [3] Dennis Brown, Yohan Baillot, Simon J. Julier, David Armoza, Joshua Eliason, Mark A. Livingston, Lawrence J. Rosenblum, and Pat Garrity. Data distribution for mobile augmented reality in simulation and training. In *Interservice/Industry Training, Simulation, & Education Conference (I/ITSEC)*, December 2003.
- [4] Dennis Brown, Yohan Baillot, Simon J. Julier, Paul Maassel, David Armoza, Mark A. Livingston, and Lawrence J. Rosenblum. Building a mobile augmented reality system for embedded training: Lessons learned. In *Interservice/Industry Training, Simulation, & Education Conference (I/ITSEC)*, December 2004.

- [5] Dennis G. Brown, Yohan Baillot, K.C. Pflueger, Simon J. Julier, and Mark A. Livingston. Virtual targets for the real world. *NRL Review*, 2005.
- [6] Dan Curtis, David Mizell, Peter Gruenbaum, and Adam Janin. Several devils in the details: Making an AR application work in the airplane factory. In *Augmented Reality: Placing Artificial Objects in Real Scenes (International Workshop on Augmented Reality)*, pages 47–60, November 1998.
- [7] Frank Dean, Pat Garrity, and Christopher Stapleton. Mixed reality: A tool for integrating live, virtual and constructive domains to support training transformation. In *Interservice/Industry Training, Simulation, & Education Conference (I/ITSEC)*, December 2004.
- [8] H.Q. Dinh, N. Walker, C. Song, A. Kobayashi, and Larry F. Hodges. Evaluating the importance of multi-sensory input on memory and the sense of presence in virtual environments. In *IEEE Virtual Reality*, March 1999.
- [9] Stephen R. Ellis and Brian M. Menges. Localization of object in the near visual field. *Human Factors*, 40(3):415–431, September 1998.
- [10] Steven Feiner, Blair MacIntyre, and Dorée Seligmann. Knowledge-based augmented reality. *Communications of the ACM*, 36(7):52–62, July 1993.
- [11] H. Fouad, James A. Ballas, and Derek Brock. An extensible toolkit for creating virtual sonic environments. In *International Conference on Auditory Displays*, 2000.
- [12] Henry Fuchs, Mark A. Livingston, Ramesh Raskar, D'nardo Colucci, Kurtis Keller, Andrei State, Jessica R. Crawford, Paul Rademacher, Samuel H. Drake, and Anthony A. Meyer MD. Augmented reality visualization for laparoscopic surgery. In *First International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 934–943, October 1998.
- [13] Chris Furmanski, Ronald Azuma, and Mike Daily. Augmented-reality visualizations guided by cognition: Perceptual heuristics for combining visible and obscured information. In *IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR 2002)*, pages 215–224, September 2002.
- [14] Joseph L. Gabbard, J. Edward Swan II, Deborah Hix, Marco Lanzagorta, Mark A. Livingston, Dennis Brown, and Simon Julier. Usability engineering: Domain analysis activities for augmented reality systems. In *SPIE (International Society for Optical Engineering)*, The Engineering Reality of Virtual Reality 2002, January 2002.
- [15] Brian F. Goldiez. *Techniques for Assessing and Improving Performance in Navigation and Wayfinding using Mobile Augmented Reality*. PhD thesis, Univ. of Central Florida, 2004.
- [16] Kiyoshi Kiyokawa, Yoshinori Kurata, and Hiroyuki Ohno. An optical see-through display for mutual occlusion of real and virtual environments. In *IEEE and ACM International Symposium on Augmented Reality*, pages 60–67, October 2000.
- [17] Mark A. Livingston, Joseph L. Gabbard J. Edward Swan II, Tobias H. Höllerer, Deborah Hix, Simon J. Julier, Yohan Baillot, and Dennis Brown. Resolving multiple occluded layers in augmented reality. In *International Symposium on Mixed and Augmented Reality (ISMAR2003)*, pages 56–65. IEEE, October 2003.
- [18] Mark A. Livingston, Lawrence J. Rosenblum, Simon J. Julier, Dennis Brown, Yohan Baillot, J. Edward Swan II, Joseph L. Gabbard, and Deborah Hix. An augmented reality system for military operations in urban terrain. In *Interservice/Industry Training, Simulation, and Education Conference*, page 89, December 2002.
- [19] Jeffrey W. McCandless, Stephen R. Ellis, and Bernard D. Adelstein. Localization of a time-delayed, monocular virtual object superimposed on a real environment. *Presence: Teleoperators and Virtual Environments*, 9(1):15–24, February 2000.
- [20] Microsoft Corporation. Directsound, 2004. Retrieved 23 June 2004 from <http://msdn.microsoft.com/>.
- [21] Ulrich Neumann and Anthony Majoros. Cognitive, performance, and systems issues for augmented reality applications in manufacturing and maintenance. In *IEEE Virtual Reality Annual International Symposium*, pages 4–11, 1998.
- [22] Wayne Piekarski and Bruce H. Thomas. ARQuake: The outdoor augmented reality gaming system. *Communications of the ACM*, 45(1):36–38, 2002.
- [23] R. Bruce Ridley. Augmented vision for automotive maintenance and repair: The benefits of testing and repair data at point of task (unpublished), Microvision, Inc., 2003.
- [24] J. P. Rolland, C. Meyer, K. Arthur, and E. Rinalducci. Method of adjustments versus method of constant stimuli in the quantification of accuracy and precision of rendered depth in head-mounted displays. *Presence: Teleoperators and Virtual Environments*, 11(6):610–625, December 2002.
- [25] Jannick P. Rolland and Henry Fuchs. Optical versus video see-through head-mounted displays in medical visualization. *Presence: Teleoperators and Virtual Environments*, 9(3):287–309, June 2000.
- [26] Michael Rosenthal, Andrei State, Joohi Lee, Gentaro Hirota, Jeremy Ackerman, Kurtis Keller, Etta D. Pisano MD, Michael Jiroutek, Keith Muller, and Henry Fuchs. Augmented reality guidance for needle biopsies: A randomized, controlled trial in phantoms. In *Medical Image Computing and Computer-Assisted Interventions*, pages 240–248, October 2001.
- [27] Andrei State, David T. Chen, Chris Tector, Andrew Brandt, Hong Chen, Ryutarou Ohbuchi, Mike Bajura, and Henry Fuchs. Case study: Observing a volume-rendered fetus within a pregnant patient. In *IEEE Visualization '94*, pages 364–368, 1994.
- [28] Andrei State, Mark A. Livingston, Gentaro Hirota, William F. Garrett, Mary C. Whitton, Etta D. Pisano MD, and Henry Fuchs. Technologies for augmented reality systems: Realizing ultrasound-guided needle biopsies. In *SIGGRAPH 96 Conference Proceedings*, pages 439–446. ACM SIGGRAPH, Addison Wesley, August 1996.
- [29] Natsuki Sugano, Hirokazu Kato, and Keihachiro Tachibana. the effects of shadow representation of virtual objects in augmented reality. In *IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR 2003)*, pages 76–83, October 2003.
- [30] Arthur Tang, Charles Owen, Frank Biocca, and Weimin Mou. Comparative effective of augmented reality in object assembly. In *Conference on Human Factors in Computing Systems (CHI 2003)*, pages 73–80, April 2003.
- [31] Anthony Webster, Steven Feiner, Blair MacIntyre, William Massie, and Theodore Krueger. Augmented reality in architectural construction, inspection, and renovation. In *Third ASCE Congress for Computing in Civil Engineering*, June 1996.
- [32] Michael White, Christopher Arena, Roy Newton, and Jeffrey Hopper. VIRTE: A training system for the future... today! In *Interservice/Industry Training, Simulation, & Education Conference (I/ITSEC)*, December 2003.